WHAT IS CLAIMED IS:

1.   In a mobile ad hoc network having a plurality of nodes, at least one node of which being a resource provider, a quality of service (QoS) aware resource discovery method, comprising the steps of:

generating at least one discovery agent from the plurality of nodes;

forming at least one dynamic domain within the ad hoc network, each dynamic domain including at least a subset of the nodes as members thereof and one discovery agent, the one discovery agent serving as a home discovery agent for its associated dynamic domain;

registering a resource by the resource provider with the home discovery agent of the resource provider's associated dynamic domain;

generating a query to discover the resource; and

discovering the resource.

2.   The method of claim 1, wherein the step of generating at least one discovery agent comprises the steps of:

broadcasting, by all eligible nodes, existence information including a node address;

electing the node that has the smallest node address as an initial discovery agent.

3.   The method of claim 2, further comprising the steps of:

selecting, by the initial discovery agent, M-1 nodes to be discovery agents, where M is a preselected number of discovery agents; and

assigning each of the selected nodes an index from the set $\{2, 3, \ldots, M\}$.

4.   The method of claim 1, wherein the step of forming at least one dynamic domain within the ad hoc network comprises the steps of:

broadcasting, by a discovery agent, a formation announcement containing at least a distance field;

comparing, by a receiving node, the distance field to a distance from the node's current home discovery agent; and

when the distance field is less than the distance from the node's current discovery agent, and when the receiving node does not have a current home discovery agent, setting the discovery agent that broadcast the formation announcement as the node's new home discovery agent.

5. The method of claim 4, further comprising the step of forwarding the formation announcement to the node's neighbors after the step of setting the discovery agent that broadcast the formation announcement as the node's new home discovery agent.

6. The method of claim 1, wherein the step of registering a resource by the resource provider with the home discovery agent of the resource provider's associated dynamic domain comprises the steps of:

sending a resource registration request to the resource provider's home discovery agent, the registration request including an attribute α of the resource;

calculating, by the home discovery agent, a hashing index β of the resource as $\beta=H(\alpha)$ in the set $\{1, 2, \ldots M\}$; and

distribute the resource registration request to discovery agents having an index of $\beta, \beta+1, \ldots, \beta+K-1$ for registration of the resource thereby, where K is a predetermined number of replications for the resource information.

7. The method of claim 6, wherein the nodes are homogeneous with a failure probability p, and wherein K is set to $|\log_{1-p}A|$, where A is a predetermined availability requirement for directory information.

8. The method of claim 1, further comprising the steps of:
discovering a failed discovery agent;
broadcasting a discovery agent selection message;
receiving responses from non-discovery agent nodes;
comparing a distance from each non-discovery agent node from which a response was received to the failed discovery agent to determine which is closest to the failed discovery agent; and

selecting the closest of the non-discovery agent nodes to be a new discovery agent for the failed discovery agent.

9.   The method of claim 8, further comprising the steps of:

assuming, by the new discovery agent, the index i of the failed discovery agent; and

recovering the failed discovery agent's directory information from at least one discovery agent having an index of i-K+1, i-K+2, . . .,i+K-1.

10.   The method of claim 1, further comprising the step of collecting, by a resource provider's home discovery agent, application-level quality of service (QoS) information from the resource provider.

11.   The method of claim 1, further comprising the step of estimating, by each home discovery agent, path quality of service (QoS) between nodes in its associated dynamic domain and those of another dynamic domain as the path latency between itself and the home discovery agent for the other dynamic domain.

12.   The method of claim 1, wherein the step of discovering the resource comprises the steps of:

transmitting, by a node to its home discovery agent, a directory query for a resource;

searching, by the node's home discovery agent, for the requested resource directory information;

when the requested resource directory information is not found, calculating, by the home discovery agent, the hashing index of the resource to determine a qualified set of discovery agents that should have the requested resource directory information;

forwarding, by the home discovery agent, the directory query to one of the discovery agents in the qualified set;

receiving, by the home discovery agent, the requested resource directory information; and

forwarding the requested resource directory information to the node.

13.   The method of claim 12, wherein the step of forwarding the directory query comprises the steps of:

determining which of the qualified set of discovery agents is closest to the home discovery agent;

forwarding the directory query to the closest discovery agent in the qualified set; and

when the closest discovery agent is unable to provide the requested resource directory information, forwarding the directory query to the next closest discovery agent until the information is found.

14.   The method of claim 1, wherein the step of discovering the resource comprises the steps of:

transmitting, by a node to its home discovery agent, a quality of service (QoS) query for a resource to be accessed;

searching, by the node's home discovery agent, for the requested resource;

when the requested resource is not found, calculating, by the home discovery agent, the hashing index of the resource to determine a qualified set of discovery agents that should have the requested resource;

forwarding, by the home discovery agent, the QoS query to at least one of the discovery agents in the qualified set;

forwarding, by the at least one discovery agent in the qualified set, the QoS query to the home discovery agents for resource providers having the resource requested;

returning, by the home discovery agents for resource providers having the resource requested, QoS and address information for the resource providers in their home domains having the resource requested;

selecting a resource provider having the best QoS; and

forwarding the address information of the resource provider having the best QoS to the node.

15.   The method of claim 14, wherein the step of forwarding, by the at least one discovery agent in the qualified set, the QoS query to the home discovery agents for resource providers having the resource requested comprises the step of multicasting.

16.   The method of claim 14, wherein the step of forwarding, by the at least one discovery agent in the qualified set, the QoS query to the home discovery agents for resource providers having the resource requested comprises the step of multiple-unicasting.

17.   In a mobile ad hoc network having a plurality of nodes arranged in dynamic domains, each dynamic domain having a home discovery agent, at least a subset of the dynamic domains having nodes therein providing resources to the network, a method of registering the resources for discovery comprising the steps of:

receiving a registration request for a resource from a node, the resource having an attribute α;

calculating a hash index of the resource as $\beta = H(\alpha)$, where β is in the set $\{1, 2, \ldots, M\}$, M being the number of discovery agents;

forwarding the registration request for the resource to discovery agents having an index of β, $\beta+1$, $\ldots$, $\beta+K-1$, where K is a predetermined number of replications for the resource registration.

18.   The method of claim 17, wherein the nodes have a failure probability p, and wherein K is set to $|\log_{1-p} A|$, where A is a predetermined availability requirement for resource registration.

19.   In a mobile ad hoc network having a plurality of nodes, a method of generating a discovery agent, comprising the steps of:

broadcasting, by all nodes, existence information including a node address;

electing the node that has a smallest node address as an initial discovery agent;

selecting, by the initial discovery agent, a preselected subset of the nodes to be discovery agents; and

wherein non-discovery agent nodes form dynamic domains by associating with one of the discovery agents.

20.   In a mobile ad hoc network having a plurality of nodes, a subset of the nodes assuming a role of a discovery agent, a method of forming a dynamic domain within the ad hoc network, comprising the steps of:

broadcasting, by a discovery agent, a formation announcement containing at least a distance field;

comparing, by a receiving node, the distance field to a distance from the node's current home discovery agent, if any; and

when the distance field is less than the distance from the node's current discovery agent, and when the receiving node does not have a current home discovery agent, setting the discovery agent that broadcast the formation announcement as the node's new home discovery agent.

21.   The method of claim 20, further comprising the step of forwarding, by the receiving node, the formation announcement after the step of setting the discovery agent that broadcast the formation announcement as the receiving node's new home discovery agent.

22.   In a mobile ad hoc network having a plurality of nodes arranged into a plurality of dynamic domains, each dynamic domain having a home discovery agent having resource directory information, a method comprising the steps of:

discovering a failed discovery agent;

broadcasting a discovery agent selection message;

receiving responses from non-discovery agent nodes;

comparing a distance from each non-discovery agent node from which a response was received to the failed discovery agent to determine which non-discovery agent node is closest to the failed discovery agent; and

selecting the closest of the non-discovery agent nodes to be a new discovery agent for the failed discovery agent.

23.   The method of claim 22, further comprising the steps of:

assuming, by the new discovery agent, the index i of the failed discovery agent; and

recovering the failed discovery agent's resource directory information from at least one discovery agent having an index selected from the group i-K+1, i-K+2, . . ., i+K-1.

24.   In a mobile ad hoc network having a plurality of nodes arranged in dynamic domains, each dynamic domain having a home discovery agent, at least a subset of the dynamic domains having nodes therein providing resources to the network, a method comprising the step of collecting, by a resource provider's home discovery agent, application-level quality of service (QoS) information from the resource provider.

25.   In a mobile ad hoc network having a plurality of nodes arranged in dynamic domains, each dynamic domain having a home discovery agent, at least a subset of the dynamic domains having nodes therein providing resources to the network, a method comprising the step of estimating, by each home discovery agent, path quality of service (QoS) between nodes in its associated dynamic domain and those of another dynamic domain as the path latency between itself and the home discovery agent for the other dynamic domain.

26.   In a mobile ad hoc network having a plurality of nodes arranged in dynamic domains, each dynamic domain having a home discovery agent, at least a subset of the dynamic domains having nodes therein providing resources to the network, a method of discovering the nodes providing the resource comprises the steps of:

receiving a directory query for a resource;

searching for the requested resource directory information;

when the requested resource directory information is not found, calculating the hashing index of the resource to determine a qualified set of discovery agents that should have the requested resource directory information;

forwarding the directory query to one of the discovery agents in the qualified set;

receiving the requested resource directory information; and

forwarding the requested resource directory information to the node.

27.   The method of claim 26, wherein the step of forwarding the directory query comprises the steps of:

determining which of the qualified set of discovery agents is closest;

forwarding the directory query to the closest discovery agent in the qualified set; and

when the closest discovery agent is unable to provide the requested resource directory information, forwarding the directory query to the next closest discovery agent until the information is found.

28.   In a mobile ad hoc network having a plurality of nodes arranged in dynamic domains, each dynamic domain having a home discovery agent, at least a subset of the dynamic domains having nodes therein providing resources to the network, a method of discovering the resource comprises the steps of:

receiving a quality of service (QoS) query for a resource to be accessed;

searching for the requested resource;

when the requested resource is not found, calculating the hashing index of the resource to determine a qualified set of discovery agents that should have the requested resource;

forwarding the QoS query to at least one of the discovery agents in the qualified set;

receiving from the home discovery agents for resource providers having the resource requested QoS and address information for the resource providers in their home domains having the resource requested;

selecting a resource provider having the best QoS; and

forwarding the address information of the resource provider having the best QoS to the node.